AF
IFW

920476-904756

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the application of	:	A G Tomlins et al
Serial No.	:	09/470,629
Filed	:	December 22, 1999
For	:	Communications Network
Examiner	:	I P Mehra
Art Unit	:	2666
Customer number	:	23644

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to "Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450," on September 13, 2004

Name of person signing Michelle Lazowski

Signature Michelle Lazowski

APPEAL BRIEF

Honorable Director of Patents and Trademarks
PO Box 1450
Alexandria, VA 22313-1450

Dear Sir,

This appeal is from the Examiner's final rejection of 24 March 2004, confirmed in the advisory action of 24 June 2004, in which all pending claims were rejected. A timely Notice of Appeal was filed on July 13, 2004 with the required fee of $330.00.

This brief is being filed in triplicate, along with the required $330.00 fee pursuant to 37 C. F. R. § 1.17(c).

09/17/2004 SSESHE1 00000059 09470629
01 FC:1402 330.00 OP

(1) Real Party in Interest

This application is assigned to Nortel Networks Limited, who is the real party in interest.

(2) Related Appeals and Interferences

There are no related appeals or interferences.

(3) Status of Claims

This application was filed with claims 1 to 12. In the response of 23 June 2003, claims 13 to 15 were added, claims 2, 4, 5 and 6 were amended, claims 7 and 8 were cancelled, and claims 9, 10, 11 and 12 were amended. This left claims 2 and 9 as independent claims. In the response of 2 January 2004, amendments were made to claims 2, 4, 5, 9, 14 and 15. No amendments were made in the third response.

All the pending claims have been finally rejected by the Examiner, and the rejection of these claims in the office action of 24 March 2004 is appealed. All the pending claims 2-6 and 9-15 as amended during the prosecution of the application are set forth in the Appendix.

(4) Status of Amendments

A response, but no amendment, was made in response to the final rejection of 24 March 2004, and so the claims now appealed are the same as those rejected.

(5) Summary of Invention

The invention is concerned with adapting synchronous time division multiplexed (TDM) traffic at an interface between a frame based synchronous network and a cell based asynchronous network (e.g. ATM). In particular it is concerned with reducing delays at such interfaces which could be incompatible with real time traffic. The two principal causes of these delays, (frame alignment delay and adaptation layer delay) will now be explained.

Prior art interfaces would involve a frame alignment operation. Frame alignment converts the timing information of each timeslot into a particular position. In transmitting across the asynchronous network, the position allows the timeslot to be correctly identified at the far side of the asynchronous network for reuse in the TDM network. Where an asynchronous network is used to substitute TDM switches, the timeslots can be established from the positions to enable timeslot switching and/or multiplexing of these timeslots into different asynchronous ATM Virtual Circuits. This alignment operation would involve a variable delay of between zero and one frame time, on average half a frame's delay or 62.5us for SONET or SDH. This is a significant part of the 250us delay budget allocated to network switches and therefore undesirable for real time traffic at least.

A second cause of delay in practice in such interfaces is caused by processing required by adaptation layers such as AAL1 and AAL2, which involves accumulating the payload of each ATM cell from a number of frames before sending. This assembly delay could vary between a worst case (e.g. one timeslot from each of 48 frames) delay of 48 x 125 μsec, or a best case (if all the time slots from one interface are used) delay of one and a half frames (1.5 x 125 μsec) for an E1 interface.

The invention, as claimed in independent method claim 2 and corresponding independent apparatus claim 9, addresses these two concerns with two distinctive

features. Firstly there is the step of "generating pointers identifying phase of the TDM traffic" and "wherein said pointers are mapped into one or more separate asynchronous cells for transport ahead of said traffic cells." By sending the phase, it is no longer necessary to carry out frame alignment at every switch or every interface between networks. Thus the frame alignment steps can be reduced or eliminated, and the associated delays can be reduced or eliminated. This phase of the TDM traffic should not be confused with prior art, in which pointers indicate a position in the (ATM cell) payload of the TDM frame boundary which has previously been aligned to the nodal frame boundary of the transmitting equipment. In other words, where the prior art shows a position pointer with a value e.g. "3", this means the first byte of the frame is positioned 3 bytes from the start of the payload of an asynchronous cell. In contrast, a pointer indicating phase and having a value e.g. "29" could mean the start of the asynchronous cell payload is the 29^{th} byte or timeslot of a synchronous frame. This value is therefore relative to a reference TDM time, not a positional value relative to an asynchronous cell payload structure.

The prior art positional pointer requires that frames having different phases (e.g. several individual E1s) must first be aligned to have a common start time (the nodal frame boundary) before they can be transmitted. The delay caused by this is avoided by using phase pointers, since the different phases can be maintained rather than removed, and simply indicated by the pointers. Another advantage is that the phase pointers can be altered easily to account for the interface having a clock frequency different to the nodal frequency. This would be achieved by making a frame one byte larger or smaller than usual, effected by changing the phase without any consequent change in the cell payload structure. This should be compared with the prior art that does not adjust cell payload position pointers to account for any change of information in the TDM interface, and must therefore slip a whole frame for such frequency variations.

Secondly, there is the step of "mapping said synchronous frames into primary multiplexed groups, mapping each said primary multiplexed group into traffic cells in a respective asynchronous virtual circuit". This multiplexing of frames before mapping into asynchronous cells, not only necessarily preserves the frame context to which the phase information refers, it means that cells can be filled and sent more quickly. The prior art does not show this. In AAL-1 and AAL-2 for example there may be one or more timeslots from many successive frames, but these frames are not multiplexed since they are not contemporaneous. Wherever the frame context of an interface is being preserved, the cells are intended to be the contents of successive frames from one interface. Hence the multiplexing reduces the waiting for enough frames to fill up cell payloads.

(6) Issues

There are two issues, firstly the rejection of claims 2-6 and 9-13 under 35 U.S.C.§103 as being obvious over Won and Irwin in view of Caves and other prior art concerning certain dependent claims, and secondly, the rejection of claims 14 and 15 under 35 U.S. §112, ¶1, as not being supported by the specification.

(7) Grouping of Claims

The claims may be considered as a group.

(8) Argument

8.1.1 35 U.S.C. §103 rejection:

Won (US6,510,163) is cited in item 5 of the final office action of 24 March 2004 as showing an interface between an asynchronous and a synchronous network, but not showing the first or second distinctive feature set out above. Won uses the

adaptation layer AAL-1. AAL-1 uses a pointer when there is a structure boundary occurring within the next two cells. The structure boundary is the boundary of the frame which has been aligned to the reference of the transmitting equipment. In general it has an arbitrary position relative to the cell payload start position. The pointer specifies the offset between these positions. It does not directly identify timeslot identity or the phase relative to a reference TDM time. Furthermore in AAL-1 SDT the bytes are of timeslots that have all been frame aligned, and it does not have the capabilities to multiplex the frames of more than one interface by interleaving. This all requires delay at the outset, and frame phase information is lost, as explained above. Hence this reference does not lead a skilled person towards the two distinctive features of the invention, nor their advantages of enabling fewer frame alignments and reduced adaptation layer delays.

Irwin (US 5,862,136) is cited in item 5 of the final office action of 24 March 2004 as showing an interface between synchronous frames and asynchronous cells. It is also cited as showing pointers and the two distinctive features of the present claims, other than the pointer identifying phase of the TDM traffic. This is respectfully traversed as will be explained below.

Caves (US 6,266,343) is cited as showing a pointer identifying phase of TDM traffic. This is again traversed as will be explained below.

Finally the Examiner tries to argue that the combination of these three references would lead to the invention because it would be obvious "to use the distinct identifier " (of Won?) "and the adaptation process of mapping pointers and labels into cells for onward transmission across the ATM network as taught by Irwin and Caves." This is respectfully traversed. The Examiner also tries to argue that: "The motivation to do so would have been to reconstruct or restore the frames in TDM format at the end of the ATM network". Again this is respectfully traversed as will be explained below.

8.1.2 Irwin:

This reference describes combining ATM and STM elements within a switching fabric. Multiple 64kb/s connections are arranged in a single cell so that more efficient use can be made of an ATM switching fabric with reduced packetisation delays. There is mention of segmenting a TDM frame into asynchronous cells, and including a distinct identifier in each cell header to enable rebuilding of the frames. This identifier is therefore only an indication of sequential order, in case the cells are lost entirely or lose sequence in asynchronous transmission. This identifier is not the phase of the frame relative to any time reference, and therefore this is not identifying phase of TDM traffic.

There is discussion at col 14 line 45 onwards that "the ATM switch removes any fixed timing relationship between a cell arrival and the synchronous 8 KHz rate used to interconnect isochronous services. To restore the relationship of channels to the 8 kHz frame structure, a block of memory, exemplified as the DS0 TSIC 420 may be used to reassemble the TDM frame structure." This reassembly of the TDM frame structure as well as the assembly of the cells themselves relies on conventional frame alignment, which has its associated half frame average delay as discussed above in the summary of invention section. Therefore Irwin teaches that such delay is inevitable and suggests no way of avoiding it. Certainly there is no suggestion of sending phase, or of sending it in separate packets, or of the advantage that the number of frame alignment operations and thus half frame time delays, can be reduced.

8.1.3 Caves:

This reference discusses sending TDM traffic from a narrowband network over a number of ATM networks to another TDM narrowband network. There is discussion of using adaptation layers AAL-1 and AAL-2. AAL-1 pointers have been discussed

above in regard to Won. AAL-2 uses a pointer in every cell to signify the offset to the boundary of a minicell header which can move relative to the ATM cell boundary. If the minicell is carrying a TDM structure, and this would represent only timeslots from a single E1 or T1, then implicitly this is the structure boundary. Hence the pointers signify the difference in position of the structure or frame boundary from the start position of the ATM cell payload. However the pointer does not directly identify timeslot identity or the phase relative to a reference TDM time. There is no teaching that frame alignment can be avoided, so there is no suggestion leading a skilled person towards sending phase, nor towards the advantage of the invention that the number of frame alignment operations and thus half frame time delays, can be reduced. Furthermore, there is no suggestion that the delays involved in using adaptation layers can be avoided.

8.1.4 The combination of these references:

Won, Irwin and Caves all rely on frame alignment of all interfaces to the nodal frame reference and do not suggest any way of avoiding it. Neither the distinct identifiers of Irwin which only serve to indicate the order of parts of a frame, nor the pointers of Caves or Won which only serve to indicate the position of parts of a frame, identify phase relative to a TDM reference. Irwin is concerned with structures within a switch, whereas Won and Caves are concerned with sending structures across networks of different types. Hence there is no incentive to combine them. Even if they were combined, the combination only serves to reinforce the teaching that frame alignment and its inevitable delay cannot be avoided. For these reasons the main claims are not obvious over the cited references taken alone or in combination.

8.1.5 Summary of previous arguments about Irwin

In the final office action, the Examiner repeated the rejection of the independent claims and some of the dependent claims for obviousness over Won in view of Irwin

and Caves. In the last response the applicants explained clearly that Irwin does not show or suggest the claim features of:

a)"pointers identifying phase of the TDM traffic"

b)"said pointers are mapped into one or more separate asynchronous cells for transport ahead of said traffic cells."

In the final action in item 8 the Examiner responded with two points relating to Irwin. A first point was that Irwin does show the header being separated from the payload. A second point was that the adaptation of cell headers into the header data stream of Irwin is an example of claim feature b). The applicants responded by explaining that the first of these points does not bring Irwin any closer to showing or suggesting the distinguishing claim features a) and b) or their advantages. It was also explained why the second point is not supported.

In the advisory action, the Examiner states it is admitted as obvious that TDM implies time slots. What relevance this has to the distinctive claim features a) and b) is not apparent. If the Examiner means to imply that this makes feature a) obvious, this of course does not follow logically, because feature a) is not the use of time slots in TDM per se, it is having phase identifiers in pointers for adapting the TDM traffic for an asynchronous network. Presumably the Examiner means to imply something else, and it would be useful to have further explanation.

In the advisory action, the Examiner goes on to address the issue of Irwin not disclosing or suggesting that the header stream or any other part of the received cells is put into separate asynchronous cells for transport. The Examiner cites Irwin as disclosing:

that demultiplexer 412 distributes ATM cells across lines 1-n, labelled 416, in accordance with the header information (480 fig. 5, refer to col. 13 lines 33-36).

In fig. 5, headers are separated from cell traffic (payload), refer to bus 414 in fig. 5, for transport ahead of payload separately to be reunited by controller 490 at bus 413c, refer to col. 13 lines 28-35.

Neither of these parts of Irwin shows or suggests the distinctive claim features a) and b), and the Examiner presents no explanation of the relevance of these passages. The first part says that ATM cells are distributed across output lines according to header information. This is in the context of a function of receiving ATM cells, carrying out octet interchange, then sending altered ATM cells. The Examiner has not explained how this is of any relevance to phase, or to synchronisation which is enabled by sending phase information. The selection of which of the output lines 1-n of Irwin labelled 416 to use is presumably determined based on destination information in the header, not on the identifier in the header. There is evidence for this by reference to corresponding part 316 described at lines 32-37 of col 12 which refers to incoming cells "destined for a particular one of the communications lines 316". Hence there is no suggestion in this passage of distinctive feature a) of a pointer identifying phase. Nor is there any suggestion of distinctive feature b) of pointers being mapped into one or more separate asynchronous cells for transport.

The second passage cited by the Examiner refers to buses carrying headers and cell structured payloads which are merged to provide a high speed stream of multiplexed cells to a demultiplexer for sending out as asynchronous cells or converting to synchronous frames. Again this does not suggest the distinctive claim features, as will be explained in more detail below.

8.1.6 Feature a), phase identifier

It seems the nearest that Irwin comes to feature a) is an identifier in each cell header:

"...the TDM frame be segmented into one or more cells..... A distinct identifier in each cell header is used so that the frame structure may be randomly reconstructed..." (col 13 lines 14-18).

As the purpose of the identifier is reconstruction rather than synchronisation of the reconstructed frame, the identifier must be cell ordering information, not phase information. This is confirmed by the discussion of frame realignment in Irwin as mentioned above, and the fact that Irwin says there is an identifier in "each" cell header, whereas phase information need not be provided in "each" cell header.

Since no further explanation or argument is presented, the Examiner seems to effectively acknowledge there is no disclosure in Irwin which suggests the identifier could be phase information, as per claim feature a), rather than conventional cell order information.

Caves is also alleged to show pointers identifying phase. This reference only shows pointers which indicate the position of the start of a frame relative to a header in a cell. This points to a location in the cell, to enable "the data structure to be recovered at the receiver". This is not a phase identifier, since it is a pointer relative to an asynchronous cell structure and so is not timing information and cannot be used for re-synchronising a time offset.

8.1.7 Feature b), pointers into separate asynchronous cells for transport ahead of traffic cells

Regarding feature b), the closest Irwin seems to come is in showing that for incoming ATM cells the cell headers are separated from cell payloads. The headers presumably contain the cell ordering identifier. However the identifiers are not mapped into separate cells as claimed. Instead, the separated headers are fed into a "header data stream" used for reconstructing a temporal image of the TDM frames,

to enable buffering and reordering of octets for onward transmission by a synchronous or an ATM link. Outgoing headers are generated by output controller 490, associated with the outgoing cell-structured payload datastream for onward transmission as ATM cells or conversion into TDM frames for onward transmission by synchronous TDM. As the header datastream is a series of headers each having an associated cell structured payload at all times, this header datastream cannot be regarded as properly separated. The header stream is presumably always synchronised with the payload stream within the switch, as no other mechanism for associating each header with its payload is shown. Certainly the headers are not put into "separate asynchronous cells" as claimed. Instead in Irwin, each header is synchronised with its payload for transport as an un-separated cell. Hence claim feature b) is effectively exactly the opposite of what is shown in Irwin.

Relating to sending pointers, Irwin teaches separating the headers within the switch only, where synchronisation between header and payload of each cell can be maintained. This is exactly opposite to the claimed feature of sending the pointers in asynchronous cells ahead of the traffic. For transport, Irwin teaches reuniting the pointers. This is exactly the opposite to claim feature b).

8.2 Claim rejection 35 USC § 112

Concerning items 2, 3 and 8, of the final action relating to an alleged lack of support in the specification for claims 14 and 15, these claims read as follows:

14. A method as claimed in claim 2, wherein each of the frames comprises a plurality of time slots, the interface comprises a timing reference and wherein the pointers identify first time slots in the frames relative to the timing reference.

15. An arrangement as claimed in claim 9, wherein each of the frames comprises a plurality of time slots, the arrangement comprises a timing reference and the means

for generating pointers is arranged to generate pointers identifying first time slots in the frames, relative to the timing reference.

In the May 24, 2004 response, the applicants argued that there is support for the plurality of time slots in any mention of TDM of course, and they are shown in many of the figures, e.g. fig 6 . There is support for a timing reference in any TDM system of course, and there is an example mentioned at page 8 line 21 which refers to "The phase offset to the node frame reference is indicated by a pointer".

In the May 24, 2004 response it was also argued that there is support for the feature of pointers identifying first time slots at page 11 lines 13 to 15, which state "Within each pointer byte the pointer value is transmitted in a 5 bit field and indicates the number of the first time slot of the E1/T1 transmitted in the node 125 μs frame."

The Examiner nevertheless argued in the advisory action that the whole of claims 14 and 15 (all features) are not supported by the specification, and that Claims 14 and 15 do not read on fig. 6. The Examiner also argued that the claim limitation "the interface comprises a timing reference and wherein the pointers identify first time slots in the frames relative to time reference", as recited in claims 14 and 15 is not supported by either drawings or specification, referring to fig. 6 or specification, pages 10-11.

The applicant is not asserting that these claims clearly read on to figure 6, but the timing reference is clearly shown by page 7 lines 15 to 17 which states "each incoming E1/T1will in general have its own phase relative to other E1/T1s and the equipment frame reference". This equipment frame reference can only be a timing reference, and the equipment corresponds to the claimed interface. Another example of the same thing is referred to on page 8 as the "node frame reference", and mentioned in figure 5 as the "equipment frame synchronisation reference". So there is ample support for the timing reference.

The remaining features of these claims are clearly shown by the passages of the application set out above, in particular page 11 lines 14 and 15 which states the pointer "indicates the number of the first time slot of the E1/T1 transmitted in the node 125 μs frame." In this case, the "node 125μs frame" is a shortened reference to an example of the node frame reference, which corresponds to the claimed timing reference. These claim features are also shown schematically in figure 5 at part "c) Byte Synchronised Signals" which shows a pointer value of 24 which is the number of the first time slot in the frame relative to the timing reference. The timing reference is shown by the vertical dotted line marked "equipment frame synchronisation reference".

For these reasons, all the claim features of claims 14 and 15 are clearly and fully supported in the text and figures of the present application and there is no new matter.

Conclusion

For the reasons explained above, the Examiner has been demonstrated to be in error Reversal is urged.

September 13, 2004

Respectfully submitted,

William M. Lee, Jr.
Registration No. 26,935
Barnes & Thornburg LLP
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 214-4800
(312) 759-5646 (fax)

Appendix

Pending Claims:

2. A method of adapting synchronous time division multiplexed (TDM) traffic at an interface between a synchronous network in which the traffic is transported in frames and an asynchronous network in which adapted traffic is transported in cells, the method comprising generating pointers identifying phase of the TDM traffic, mapping said synchronous frames into primary multiplexed groups, mapping each said primary multiplexed group into traffic cells in a respective asynchronous virtual circuit, and wherein said pointers are mapped into one or more separate asynchronous cells for transport ahead of said traffic cells.

3. A method as claimed in claim 2, wherein said primary multiplexed groups are multiplexed by byte interleaving into a secondary multiplexed signal.

4. A method as claimed in claim 3, wherein said asynchronous network is an ATM network carrying ATM cells and said adaptation is performed using ATM adaptation layer zero (AAL0).

5. A method as claimed in claim 4, wherein a time slot group frame boundary coincides with a start of an ATM cell.

6. A method as claimed in claim 2, wherein each cell containing time slots from a synchronous frame is given its own virtual channel indicator, and wherein cells relating to that frame are allocated a common virtual path indicator so that said cells can be transmitted and switched together.

9. An arrangement for adapting synchronous time division multiplexed (TDM) traffic at an interface between a synchronous network in which traffic is transported in synchronous frames and an asynchronous network in which adapted traffic is transported in cells, the arrangement comprising means for generating pointers

identifying phase of the TDM traffic, mapping means for mapping said synchronous frames into primary multiplexed groups and for mapping each said primary multiplexed group into traffic cells in a respective asynchronous virtual circuit, and wherein said pointers are mapped into one or more separate asynchronous cells for transport ahead of said traffic cells.

10. An arrangement as claimed in claim 9, and including means for multiplexing said primary multiplexed groups by byte interleaving into a secondary multiplexed signal.

11. An arrangement as claimed in claim 9, wherein each cell containing time slots from a synchronous frame is given its own virtual channel indicator, and wherein cells relating to that frame are allocated a common virtual path indicator so that said cells can be transmitted and switched together.

12. An arrangement as claimed in claim 9, and provided in the form of an integrated circuit.

13. An arrangement as claimed in claim 9, wherein said asynchronous network is an ATM network carrying ATM cells.

14. A method as claimed in claim 2, wherein each of the frames comprises a plurality of time slots, the interface comprises a timing reference and wherein the pointers identify first time slots in the frames relative to the timing reference.

15. An arrangement as claimed in claim 9, wherein each of the frames comprises a plurality of time slots, the arrangement comprises a timing reference and the means for generating pointers is arranged to generate pointers identifying first time slots in the frames, relative to the timing reference.